

LA-UR- 03 -1560

Approved for public release;  
distribution is unlimited.

C-1

Title: ParaView: An End-User Tool for Large Data Visualization

Author(s): James Ahrens, Berk Geveci, Charles Law

Submitted to: Visualization Handbook, Academic Press.



99

**Los Alamos**  
NATIONAL LABORATORY



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

# ParaView: An End-User Tool for Large Data Visualization

James Ahrens

Los Alamos National Laboratory

Berk Geveci, Charles Law

Kitware, Inc.

## Abstract

This paper describes the design of a visualization tool, called ParaView<sup>i</sup>, a tool for scientists to visualize and analysis extremely large data sets. The tool provides a graphical user interface for the creation and dynamic execution of visualization tasks. ParaView transparently supports the visualization and rendering of large data sets by executing these programs in parallel on shared or distributed memory machines. ParaView supports hardware-accelerated parallel rendering and achieves additional rendering performance via level-of-detail techniques. Contributions of this architecture are a design that offers the ease of using a graphical user interface to interactively analyze and visualize data along with an automatic means for using these user interface interactions to create a visualization program. The design balances and integrates a number of diverse requirements including the ability to handle large data, ease of use and extensibility by developers. This paper describes these requirements, identifies their importance to scientific users, and discusses key design decision and tradeoffs.

## Introduction

Sensors and scientific simulations are generating data at unprecedented rates making data visualization with traditional visualization solutions difficult or even impossible. In order to address the simulation scientists' visualization needs we spoke with simulation scientists and gathered a set of requirements. The high-level requirements that guided the design of ParaView are support for an efficient and effective workflow and support for the visualization and analysis of extremely large datasets. The challenge was to create a design and implementation that meet both these complex requirements and balanced conflicts between them.

## Workflow requirements

Visualization is one task of many for simulation scientists, other simulation tasks include, theoretical work, programming, problem setup, analysis and data management. Therefore, the first workflow requirement is efficient visualization tool ease of use, that is, how long it takes to create results and what visualization domain knowledge is required to run the tool, determines whether the tool is used and how often. A coarse approximation of the simulation scientists' workflow includes two modes, an exploratory mode, in which an interactive GUI-based tool is used to explore a dataset, and a batch mode, in which a scripting or programming language is used to write and execute a program that creates an animation. Thus the second workflow requirement is equal support for both modes. This coarse approximation can be refined further identifying how data is input (during the simulation run or after processing of the simulation is complete) and what type of interface is used (GUI, scripting, VR)<sup>1</sup>. Addition workflow requirements include tool portability, accessibility and extensibility. Portability is required because of the diverse collection of resources available to scientists to run their simulations and visualizations. Tool accessibility is the ability to quickly gain access to, setup, modify if necessary and run the tool. Open-source projects are more accessible since the package is

---

<sup>i</sup> ParaView is an open-source package and can be downloaded at [www.paraview.org](http://www.paraview.org).

typically available on the Internet and any localization issues can quickly be address because the source is available. We define extensibility as the ability to easily add new functions and graphical interfaces to the tool.

## **Large data visualization requirements**

The ability to handle large data is also a critical requirement. We define large data as data that exceeds the resource limits (i.e. the elements of the storage hierarchy – memory, disk, tape) of a single machine. The first aspect of the large data handing requirement is a functional one; can the data be visualized at all? Techniques such as data streaming (i.e. processing the data incrementally) and parallelism can be used to process large data sets. Workflow requirements, such as portability mandate that these are available on both shared and distributed-memory parallel machines. The second aspect of the large data handling requirements is performance; can the data be processed quickly? Techniques such as multi-resolution representations and parallelism can be used to improve both visualization and rendering performance.

## **Related Work**

There are a number of visualization packages available for use by scientists. Each of these packages meets a subset of the identified requirements. In this section, we will discuss a few of these packages, specifically AVS<sup>2</sup>, OpenDX<sup>3</sup>, SCIRun<sup>4</sup>, Ensight<sup>5</sup>, identifying their strengths and describing which requirements they meet. ParaView was designed to meet all workflow and large data and visualization requirements.

## **Workflow requirements**

Ensignt, ParaView and are turnkey/GUI-based visualization applications and a graphical user interface to create visualization tasks. AVS, OpenDX, SCIRun use dataflow program graph editor to compose programs. Data-flow program graph editors were thought to provide a good tradeoff between the needs of visualization developers and end-users: for developers by providing the ability to create complex program graphs and for end-users by providing a graphical interface to create these graphs. In practice, learning visual programming with data flow graphs is considered by many scientists a significant barrier to creating visualization tasks and thus they consider GUI-based interfaces easier to use. OpenDX, SCIRun and ParaView are all open-source packages making them easily accessible and extensible. All packages offer interactive and batch interaction modes. SCIRun provides support for computational steering, interacting with and visualizing data from a running simulation. In contrast with other packages, ParaView uses a general purpose scripting language, Tcl, for batch commands. Advantages of using a general purpose scripting language are general purpose computing functionality is available for use and implementation of the language is supported by others.

## **Large data visualization requirements**

All these packages are portable to most architectures when run on a single machine. Differences arise on their portability to parallel architectures. AVS, OpenDX, SCIRun all support parallel execution on shared memory machines. They also all rely on a centralized executive to allocate memory and execute programs. This reliance makes it difficult to port these packages to distributed memory machines. Ensight uses a client/server architecture and currently provides a shared-memory implementations of both the client and server. Ensight also has a distributed-memory port of the server. ParaView is portable to both shared and distributed-memory machines. ParaView is the only listed package that can incrementally process data.<sup>ii</sup>

---

<sup>ii</sup> ParaView's data streaming feature is available in batch mode. The DDV package

The ability to process datasets bigger than the available computing resources is a key consideration when processing extremely large datasets since resource availability changes over time.

## Design

ParaView is designed as a layered architecture. The foundation is the visualization toolkit (vtk).<sup>67</sup> It provides the foundation of the tool: data representations, algorithms and a mechanism to connect these representations and algorithms together to form a working program. The second layer is the parallel extensions to the visualization toolkit. The parallel vtk layer, extended vtk to support streaming of all data types and parallel execution on shared and distributed memory machines.<sup>iii</sup> The third layer is ParaView itself. ParaView transparently supports the visualization and rendering of large data sets via hardware acceleration, parallelism and level-of-detail techniques. ParaView provides a graphical user interface to interactively analyze and visualize data along with an automatic means for using these interactions to create a visualization program. Each layer meets a subset of the requirements and adds additional functionality to the layer below.

### Visualization Toolkit

The Visualization Toolkit is the foundation of the ParaView architecture. Vtk provides data representations for a variety of grid types including structured (uniform, non-uniform rectilinear grids as well as curvilinear grids), unstructured, polygonal and image data and examples of these grid types are shown in Figure 1. Vtk provides hundreds of visualization and rendering algorithms that process these data types including isosurfacing, cutting/clipping, glyphing and streamlines. Vtk also provides algorithms for polygonal and volume rendering and a keyboard and mouse-based interaction model. Algorithms are organized into dataflow program graphs and a demand-driven dataflow is used to run these programs. Core functionality in Vtk is written in C++. In order to use the toolkit, Vtk offers both a C++ library interface and set of scripting interfaces including java, python and tcl interfaces. The library interface provides the best performance and the scripting interfaces offer the advantage of rapid prototyping of programs. The toolkit is tested automatically once a day and continuously, whenever a developer commits a change, continuous dashboards are run and submitted, using an open-source testing framework that improves the toolkit's reliability. The toolkit provides the basis for ParaView's portability, accessibility, full range of features and support for interactive and scripting usage modes. More details on VTK can be found in this book in a related chapter on the toolkit.


### Parallel and Distributed Visualization Toolkit

Additional functionality was added to the Vtk to support data streaming and parallel computation.<sup>8</sup> Both functions rely on the ability to break a dataset into smaller pieces. Data streaming incrementally processes these smaller pieces supporting the processing of an extremely large dataset with computing resources that cannot store the entire dataset at one time (either in memory or on disk). Data streaming requires that all VTK data types are separable into pieces and the toolkit's algorithms correctly process these pieces. To process pieces in a dataflow pipeline, a mapping must be defined that specifies for each algorithm what portion of the input data is required to generate a portion of the output data. With this information, algorithms can generate only pieces of their output for a given input. Each algorithm must insure program results must be invariant, regardless of the number of pieces used. These requirements were met by creating a partitioning of both structured and unstructured grids types and providing ghost levels

---


<sup>iii</sup> These extensions are currently part of the toolkit but were added after the original design of the toolkit was complete.

which are points/cells that shared between processes and are used by algorithms which require neighborhood information. VTK's structure grid types use extents to partition these data types into pieces. A piece of a structured grid is defined by its extent which describes a contiguous range of elements in each dimension (i.e. in 3D a sub-block of a complete block). VTK's unstructured grid types use an element of a collection scheme (i.e. piece M of N). A procedure for converting between grid types was also defined. For example, each structured extent piece maps one-to-one to an unstructured piece. Additional policies take care of handling boundary conditions and creating ghost levels for all grid types. This data streaming ability supports data parallelism. Instead of processing pieces one of the time, in each processor processes a different piece based on its processor id. Examples of dataset partitioning and the creation of ghost levels is shown in Figure 2. Figure 1 shows a CTH non-uniform rectilinear grid data set that was processed in 8 pieces. The original data set contained cell centered attributes. Volume fraction attributes for both the projectile and plate were first interpolated to vertices before an isosurface filter was used to extract the material surfaces. Both the vertex interpolation and normal generation require ghost cells to insure partition invariant results.



**Figure 1: The figure shows the different types of data sets that Vtk and ParaView can handle. The upper left dataset is a uniform rectilinear volume of an iron potential function. The Upper right image shows an iso surface of a non-uniform rectilinear structured grid. The lower left image shows a curvilinear structured grid data set of airflow around a blunt fin. The lower right image shows an unstructured grid data set from a blow molding simulation.**

Parallel communication and control classes encapsulate details of process initialization and communication libraries such as a shared-memory implementation or MPI. The streaming and parallel computing features can be access both as a C++ library interface and set of scripting interfaces. These feature extensions provide the basis for ParaView's large data functionality and performance requirements.



**Figure 2:** The image on the left was generated with a CTH data set partitioned into eight pieces. Each piece was assigned a different color. The image on the right shows only one partition with six extra ghost levels. The cells are colored by ghost level. In practice, usually only one ghost level is necessary.

## ParaView

ParaView provides a graphical user interface for the interactive exploration of large data sets. It builds this functionality on Vtk and the parallel and distributed extensions to Vtk. An overview of the tool from a user perspective is presented first, followed by a technical description of how the tool's functionality is achieved.

### Overview

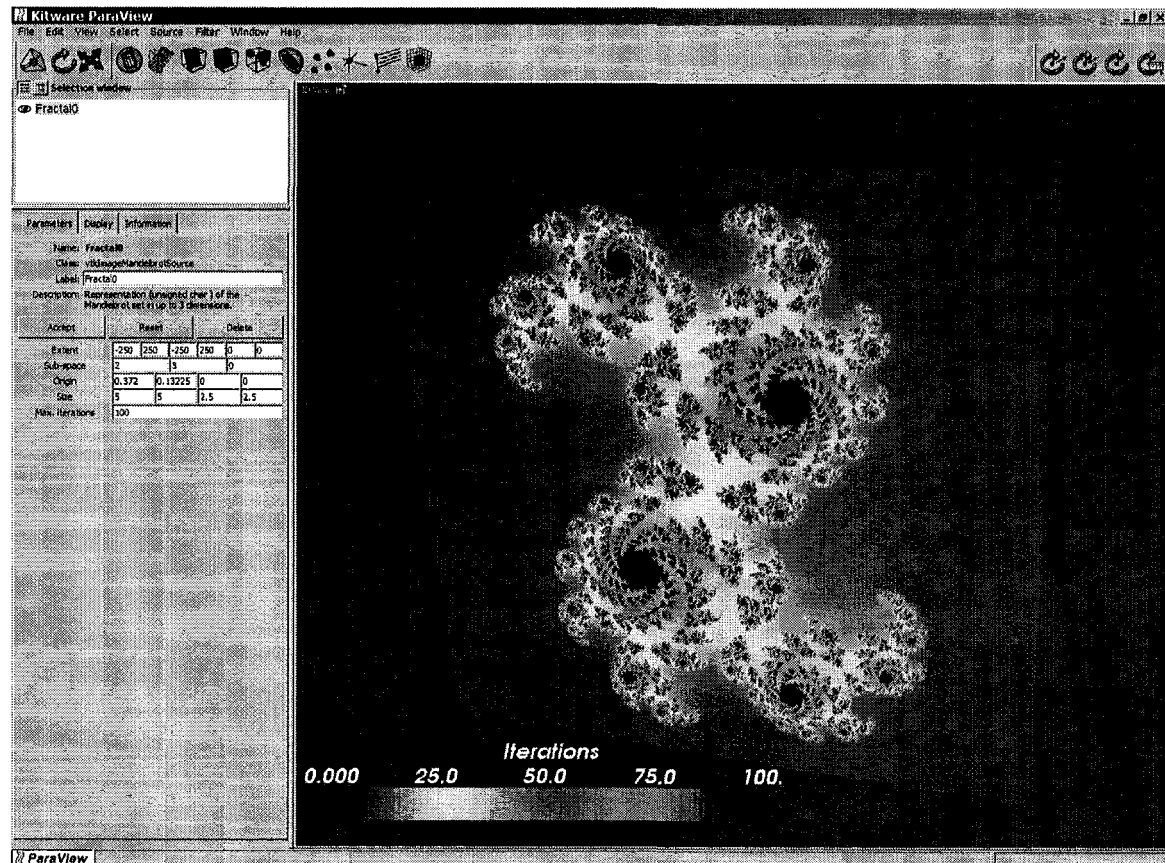
A sample ParaView session is shown in Figure 3. There are several regions to the user interface including the Menu Bar along the top of the application, the Toolbar just below the Menu Bar, the Left Panel on the left side and the Display Area on the right side. Each of these areas is described in more detail below.

**The Menu Bar:** The top menu bar provides menu buttons for loading and saving data, creating sources and filters, viewing other windows, displaying help, and other standard functionality.

**Toolbar:** The Toolbar contains buttons for resetting the camera, switching between 2D and 3D interaction modes, and changing the center of rotation. In addition, the Toolbar contains shortcut buttons to instantiate some commonly used filters.

**Left Panel:** The top portion of this panel contains the selection or navigation window. The selection window provides a list of instantiated sources and filters. The navigation window provides a dataflow program graph representation of the user's task. The area below the navigation window is where the properties of sources and filters are set, which we refer to as a property sheet, application and 3D display settings such as background color, level of detail parameters and camera controls are located, and animations can be created.

**Display Area:** The Display Area is where the 3D representation of the scene is rendered. Mouse and keyboard interaction are provided in this area.



**Figure 3: ParaView**

To add new filters the user selects a source or filter from the Source or Filter menus on the Menu Bar. Sources include various types of readers or computer-generated sources. A sample of the possible filters includes:

- *Contours and isosurfaces* can be extracted from all data types using scalars or vector components. The results can be colored by any other variable or processed further. When possible, structured data contours/isosurfaces are extracted with fast and efficient algorithms which make use of the special data layout.
- Vectors fields can be inspected by applying *glyphs* (currently arrows -*vector arrows*-, cones and spheres) to the points in a dataset. The glyphs can be scaled by scalars, vector component or vector magnitude and can be oriented using a vector field.

- A sub-region of a dataset can be extracted by *cutting* or *clipping* with an arbitrary plane (all data types), specifying a *threshold* criteria to exclude cells (all data types) and/or specifying a *VOI* (volume of interest - structured data types only)
- *Streamlines* can be generated using constant step or adaptive integrators (currently, streamlines are not available in parallel, this feature is under development). The results can be displayed as points, lines, tubes, ribbons etc. and can be processed by a multitude of filters.
- The points in a dataset can be *warped* (displaced) with scalars (given a user defined displacement vector) or with vectors (unavailable for rectilinear grids).
- With the *array calculator*, new variables can be computed using existing point or cell field arrays. A multitude of scalar and vector operations are supported.
- Data can be *probed* on a point or along a line. The results are displayed either graphically or as text and can be exported for further analysis.

ParaView provides many other data sources and filters by default (edge extraction, surface extraction, reflection, decimation, extrusion, smoothing...) and any Vtk source or filter, of which there are hundreds, can be added to ParaView by providing a simple XML description for its user interface for its property sheet.

The Source and Filter menu are dynamically updated to contain a list of sources/filters that can input the output of the currently selected module (i.e. source / filter). The selected module is either the last module created, or the one most recently selected from the Select menu or Selection/Navigation Window. Once a module is chosen from either menu, a new instantiation of the module is created and the module's property sheet is displayed. In this manner, a data-flow program graph is created. In order to manipulate or view the properties of a module, the module is selected and its property sheet is shown and the user can view or edit the listed values.

## Meeting the workflow requirements

ParaView simplifies its use by minimizing the knowledge of dataflow programming required by its users in order to use the tool. Specifically, a user can specify simple tasks, for example, creating a source and applying simple filters, without needing to be aware of dataflow programming. This is because ParaView's default behavior is to add new modules to the last module created. When the user wants change this behavior, for example to apply another filter to the source data, they can use the selection window to reset the location where the new module will be applied. ParaView also simplifies the choice of a new module by only listing modules which accept the correct data type for insertion. For advanced users who wish to create complex program graphs, the program graph is available for manipulation in Navigation window. ParaView is designed so that visualized results dominate the GUI real-estate and the manipulation of program graphs is relegated to a much smaller area. This allows scientists to focus on their visual analyze and not on visual programming which is typically of secondary importance to them.

When modules are instantiated in ParaView they create visual output in the display area providing immediate visual feedback to the user about their data and the effect of the applied module. For example, as shown in Figure 1, when the user creates the two-dimensional source, in this case, a fractal source, ParaView automatically creates a color mapping of the data. This feature improves ease of use because it provides default settings for the module parameters, freeing the user from this task. This feature does have a downside; it hampers the ability to stream data since every module instantiation would cause ParaView to stream visual results. For now, we have chosen only permit data streaming in batch mode. Solutions to this problem include:



offering the option of turning on and off interactive data streaming as well as offering the option of turning on and off the immediate feedback feature.

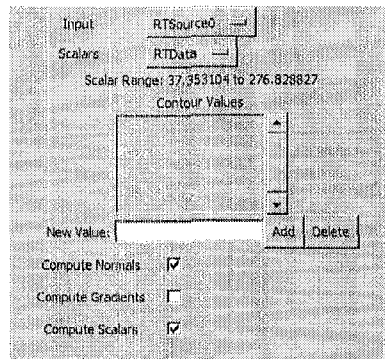
Users can change the parameters of some modules directly by interacting with the 3D view using 3D manipulators. For example, the user can manipulate the seed line of a streamtrace filter by clicking on a control point and dragging the line to the new location. There are also 3D manipulators for probing a dataset with a point or line and cutting or clipping a dataset with a sphere or plane. 3D manipulators improve ease of use by allowing users to quickly apply visualization modules to datasets by letting them interactively use the mouse to select location parameters instead of setting the parameters numerically in the user interface. When the manipulator is adjusted interactively in the 3D view, the numerical values of the location parameters are updated in the user interface and the user can then fine tune these values. For example, the 3D manipulator can be used to select a coarse location for the line probe and the user can refine these parameters to select a more exact location.

ParaView is portable to most architectures. To achieve this, only packages that work across a large number of platforms were used in developing ParaView. For example, to achieve a portable user interface, Tk, was chosen. Tk, is an interpreted scripting language and is the graphical user interface companion to the Tcl scripting language. The application framework at the core of ParaView is a unique blend of Tcl/Tk and C++. Tk is used as the widget set but C++ objects, which provide higher-level UI components, are created to encapsulate the widgets. Like Vtk objects, these C++ UI objects are automatically wrapped in Tcl.

ParaView's user interface can be modified and extended both statically with XML configuration files and dynamically, at runtime, using the Tcl scripting interface. All ParaView modules and their corresponding user interface are initialized by parsing XML based configuration files. These files contain the input/output types of the modules, if applicable, an icon name to be displayed on the toolbar, a list of widgets to be displayed on the module's parameter page, the corresponding module parameters and in case of reader modules information about the file type. For example, Figure 4 present the XML description listed in ParaView default configuration file and corresponding user interface for isoline/isosurface module.

```
<Module name="Contour" class="vtkPVContour" module_type="Filter" root_name="Contour"
  button_image="PVContourButton" output="vtkPolyData" input="vtkDataSet">
  <Filter class="vtkPVContourFilter"/>

  <InputModule id="im" label="Input" trace_name="Input" input_name="PVInput"
input_type="vtkDataSet"/>
  <ArrayMenu id="am" input_name="Input" attribute_type="Scalars"
    label="Scalars" input_menu="im" number_of_components="1"/>
  <ScalarRangeLabel array_menu="am"/>
  <ContourEntry label="Contour Values" trace_name="Contour Values"/>
  <LabeledToggle label="Compute Normals" variable="ComputeNormals"/>
  <LabeledToggle label="Compute Gradients" variable="ComputeGradients"/>
  <LabeledToggle label="Compute Scalars" variable="ComputeScalars"/>
</Module>
```



**Figure 4: XML description for the isolate/isosurface module in ParaView and corresponding user interface generated by the XML**

Since all ParaView widgets have corresponding Tcl representations, the GUI can be modified at runtime by loading scripts or typing commands at a command console. This allows the user to, for example, add new widgets, create dialog windows or load additional libraries at runtime. These features can be used to customize what modules are loaded and how they are presented to the user. A visualization developer can edit the ParaView GUI configuration file (or add new ones) and/or write custom scripts to customize ParaView for use by specific application users. For example, for the climate modeling community, a configuration file could be written to add a suite of climate analysis modules and customize the existing modules such as contouring to meet community conventions. ParaView meets the accessibility requirement since it is available in open-source form and with pre-compiled binaries users can be up and running immediately after they download the package.

Recall that, a coarse approximation of the simulation scientists' workflow includes two modes, an interactive mode, in which an interactive GUI-based tool is used to explore a dataset, and a batch mode, in which a program is executed to create an animation. ParaView supports both these modes. ParaView's interactive mode was detailed in the overview section. Every interaction with the ParaView GUI can be saved in a session file since every interaction has a corresponding Tcl command. The file can be reloaded to reproduce the session. Furthermore, since the session file is a Tcl script, it can be edited/modified and then reloaded to obtain different results. In addition to dataset exploration, the interactive mode is also used to create programs to run in batch. ParaView also supports saving the current program graph as a Vtk script. A series of queries allow the user to customize the script. A session script differs from a Vtk script because a session script saves every ParaView interaction (i.e. every interaction used to create a program graph) and a Vtk script saves only a program graph.

## Meeting the large data visualization requirements

ParaView supports large data visualization via techniques that include parallel processing, level-of-detail rendering and data streaming.

### Parallelism and data streaming

ParaView supports parallelism, either using shared-memory processes or distributed-memory processes via MPI. When ParaView is run in parallel, data is automatically read, processed and, if necessary, rendered in a data parallel manner. ParaView's parallel software architecture includes three types of processes, a client process that runs the graphical user interface and two types of server processes: root and slave processes. The client process communicates with the root server process by broadcasting commands. The root server process receives these commands and re-broadcasts them to all slave servers. A command is a script; currently, a Tcl script is used. In the

future, user may be able to select the scripting language to use, for example, Python or java. All servers run an interpreter and use it to execute received commands. This communication mechanism is used to create a copy of the same program graph on each process. The program graphs on the servers manipulate pieces of the full dataset for data parallel processing and the program graph on the client is used to store program state such as parameter settings for modules. ParaView's user interface elements update the client's program graph and the changes are sent as scripts to the root and slave servers. For example, when a user creates a filter module, such as an isosurface, a script is created that instantiates, sets parameters, and appends the isosurface module to a program graph. This script is then communicated to and interpreted by both the client and server processes. All processes use the same naming convention and thus one script works for all process's program graphs. All program graphs are initialized with a rendering module.


Communication between processes is handled internally by the modules in the program graphs, for example, all rendering modules communicate to implement a parallel rendering algorithm. In addition to algorithms written for the parallel Vtk extensions, such as computing ghost cells and parallel rendering, ParaView requires parallel information gathering algorithms in order to present global dataset size sets in the user interface.

ParaView currently supports data streaming in batch mode. When the user writes a batch script, an option prompts the user as to whether they would like to stream their data and what memory limit they are bounded by. Streaming along with data parallelism are effective techniques for processing large datasets and effectively fulfill ParaView's large data visualization requirement. We have used these techniques to efficiently and effectively isosurface and color a collection of datasets that ranged in size from 10's of gigabytes to approximately a petabyte in size<sup>8</sup>.

### **Level-of-detail and parallel rendering**

ParaView's rendering module supports both level of detail and parallel rendering techniques<sup>9</sup> in order to visualize very large datasets interactively. Interactive rendering of large datasets is an extremely challenging problem and therefore we applied a number of techniques to achieve good performance.

Level-of-detail techniques increase rendering performance at the expense of image quality. Two different LOD techniques are used in ParaView: geometric LOD and rendered image LOD. The geometric LOD technique creates a collection of reduced geometric resolution models. In general, models with less geometric elements render faster than with more elements. When the user is interacting with model (i.e. rotating, translating, zooming), a reduced resolution model is used in order to render quickly. When the interaction is complete, the full-resolution model is rendered. Figure 5 shows an example of the full and reduced resolution model. Vtk's quadric clustering algorithm is used to simplify surfaces. This algorithm preserves data attributes, so the LOD models have the same visual appearance of the original data. Timing results for the decimation algorithm are given in Table 1. Decimation can introduce significant visual artifacts in the rendered image. However, we have concluded that these artifacts are acceptable during interactive rendering. Decimation can also work well with the geometry redistribution technique discussed later in this section. The smaller decimated models can easily be collected and rendered locally.



**Figure 5: The full-resolution surface (left) has 821,495 triangles and renders in intermediate mode in 0.93 seconds on a GeForce2 GTS, and the decimated surface (right) has 35,400 triangles and renders in 0.04 seconds.**

|              | Time To Decimate |
|--------------|------------------|
| 1 Processor  | 6.25 Seconds     |
| 2 Processors | 3.515 Seconds    |
| 4 Processors | 1.813 Seconds    |

**Table 1: The full-resolution surface has 821,495 triangles, and the decimated surface had approximately 35,400 triangles.**

The rendered image LOD technique involves rendering to a small image and using pixel replication to magnify the image for final display. It is essentially lossy compression of the image that has no compression cost and minimal decompression cost. Although the resulting visual artifacts are noticeable, they do not significantly impair interaction with the data set. This technique reduces rendering times through all stages. Initial rendering times can be faster because fill areas are smaller. The times for getting image buffers, and compositing them are directly proportional to the area of the rendered image.

Using data parallelism in the renderer is also critical for high performance. The renderer is always in one of two states: an interactive state when the user is interacting with the GUI by rotating, translating or zooming geometry and a still state when the user is not interacting with the GUI. The steps of the rendering algorithm are described below. Note that ParaView's geometry and image data can be either serial or parallel data and this may change during execution. The algorithm below is applied to each geometric object to be rendered:

1. **If (Rendering state is interactive) then apply geometric and image LOD algorithm –**
  - a. The geometric LOD algorithm is applied when an estimate of the time to render the object (based on the number of points in the object) exceeds a user modifiable threshold. When the threshold is exceeded rendering occurs with a reduced resolution version of the object. If a reduced resolution version does not exist one is created in serial or parallel.
  - b. The image LOD algorithm is applied when the time to render the last frame exceeds a user modifiable threshold. Using the previous frame time as an estimate a new image size is calculated to reduce rendering time to below the threshold.
2. **If (Geometry data is parallel) then apply parallel geometry load redistribution algorithm –** The result of the LOD algorithm is geometry data. If there is parallel geometry

data, this data can be redistributed from its current location on the processes to a subset of the processes. For example, if the geometry is small enough (i.e. after it is reduced by Step 1a) it can be more efficient to collect and render the geometry on a single process. This avoids the cost of parallel image compositing (i.e. Step 4). In this future, this step will also be used to balance geometric load across processes for more efficient performance.

3. **Rendering** - The result of the redistribution algorithm is geometry data. A rendering operation then renders this geometry to create an image and depth buffer result of the image size set in Step 1b. Rendering can be serial or parallel, hardware or software-based, and occur either onscreen or offscreen.
4. **If (Image Data is Parallel) then apply parallel image compositing** - If there is parallel imagery, then this image data is composited together using the depth buffer to resolve conflicts to create a final image. ParaView currently supports a binary tree compositing with an option of using run-length encoding to losslessly compress images for speed<sup>10</sup>. With large window sizes and many processes, this communication time can be the major factor limiting rendering performance. Compositing transmission time grows linearly with render window area and scales logarithmically with the number of processes. This is why both Step 1a and 1b offers methods (collecting geometry to a single process and compositing using smaller images) to either skip or speedup this compositing step.

Notice that different paths through these steps are possible. For example, a reduced level-of-detail model can be rendered locally when the renderer is in the interactive state and the full resolution version of the same model can be rendered in a parallel when the renderer is in the still state.

## Results

This section presents visualization results generated by ParaView for several application areas.

Figure 6 shows isosurfaces of the visible woman data set. The 512x512x1734 900 MB dataset is composed of seven sections. Each section is a uniform rectilinear grid generated by a CT scan. Two isosurfaces were extracted, one for bone and one for skin. The skin isosurface was clipped in order to reveal the bone iso surface. One block of the skin was colored by process id to show the data partitioning. ParaView was run with 4 server processes in this example. This example also demonstrates ParaView's ability to process multi-block data sets. Many structured data sets divide the domain into blocks. Each block is configured to get the best resolution sampling for its domain. Since some data sets can have hundreds of blocks, it is important to group these blocks into a single entity that can be filtered in one step. ParaView has group and ungroup filters, which simplify processing of multi-block datasets.

Figure 7 shows streamlines generated by ParaView using a dataset of airflow around a delta wing. This example also shows the 3D line widget used to seed the streamline algorithm. The streamline filter propagates integration across partition boundaries, and can execute in parallel. The delta wing and the contour surface were obtained by extracting sub-grids from the original curvilinear dataset. Since the actual dataset contains only half the wing due to symmetry, a reflection filter was applied to all surfaces. Both surfaces were colored by mapping the stagnation energy through a default lookup table.

Figure 8 Global Climate Model

## **Conclusions**

This paper presents the design of ParaView, an end-user tool for large data visualization. ParaView provides a graphic user interface for visualizing large datasets using techniques that include data parallelism, level-of-detail and streaming to meet its workflow and large data visualization requirements. In the future, there are number directions to extend ParaView including incorporation of streaming algorithm in user interface and rendering support of extremely large datasets for tiled display walls. ParaView is open source tool and members of visualization community are invited to use and add new features.

## **Acknowledgments**

This work was supported in part by grants from the US Department of Energy ASCI Views program and the DOE Office of Science. We acknowledge the Advanced Computing Laboratory of the Los Alamos National Laboratory, where we performed portions of this work on its computing resources.




Figure 6: The Visible Woman dataset in ParaView





Figure 7: Delta Wing





**Figure 8:Global Climate image**

## References

- <sup>1</sup> R. Knight, J. Ahrens, P. McCormick, "Improving the Scientific Visualization Process with Multiple Usage Modalities", LAUR-001619.
- <sup>2</sup> C. Upson et al., "The Application Visualization System: A Computational Environment for Scientific Visualization," *IEEE Computer Graphics and Applications*, vol. 9, no. 4, July 1989, pp. 30-42.
- <sup>3</sup> G. Abrams and L. Trenish, "An Extended Data-Flow Architecture for Data Analysis and Visualization," *Proc. IEEE Visualization 1995*, IEEE CS Press, Los Alamitos, Calif., 1995, pp. 263-270.
- <sup>4</sup> S.G. Parker, D.M. Weinstein, and C.R. Johnson, "The SCIRun Computational Steering Software System," *Modern Software Tools in Scientific Computing*, E. Arge, A.M. Brauset, and H.P. Langtangen, eds., Birkhauser Boston, Cambridge, Mass., 1997, pp. 1-40.
- <sup>5</sup> K. Misegades "EnSight's Parallel Processing Changes the Performance Equation", <http://www.ceintl.com/products/papers.html>.
- <sup>6</sup> W. Schroeder, K. Martin and W. Lorensen, "The Design and Implementation of an Object-Oriented Toolkit For 3D Graphics and Visualization" *Proc. IEEE Visualization 1996*, IEEE CS Press, Los Alamitos, Calif., 1996, pp. 263-270.
- <sup>7</sup> W.J. Schroeder, K.M. Martin, and W.E. Lorensen, *The Visualization Toolkit An Object-Oriented Approach to 3D Graphics*, Prentice Hall, Upper Saddle River, N.J., 1996.

---

<sup>8</sup> J. Ahrens, K. Brislawn, K. Martin, B. Geveci, C. C. Law, M. Papka, "Large-Scale Data Visualization Using Parallel Data Streaming," *IEEE Computer Graphics and Applications*, vol. 21, no. 4, July/August 2001, pp. 34-41.

<sup>9</sup> S. Molnar et al., "A Sorting Classification of Parallel Rendering," *IEEE Computer Graphics and Applications*, vol. 4, no. 4, July 1994, pp. 23-31.

<sup>10</sup> J. Ahrens and J. Painter, "Efficient Sort-Last Rendering Using Compression-Based Image Compositing", *Proc. of the Second Eurographics Workshop on Parallel Graphics and Visualization*, 1998, 145-151.